nature portfolio

Peer Review File

Universal Gates from Braiding and Fusing Anyons on Quantum Hardware

Corresponding Author: Professor Ruben Verresen

Any redactions in this file are there to maintain patient confidentiality, the confidentiality of unpublished data, or to remove third-party material.

This file contains all reviewer reports in order by version, followed by all author rebuttals in order by version.

Version 0:

Reviewer comments:

Referee #1

(Remarks to the Author)
This paper describes experimental realization of the non-abelian topological order associated with the group S3. This is significant because, as a topological phase of matter, S3 is universal for quantum computation — meaning that operations that move and fuse its quasiparticles can be used to carry out a universal gate set. The experimental data presented indicates impressive fidelity with the actual topological state, both after initial state preparation and after elementary computationally relevant operations needed to comprise a universal gate set.

Overall I think this work represents a significant advance, worthy of publication in Nature. In particular, while it is not the first experimental demonstration of preparing non-abelian topological order, the demonstration of a universal set of operations is both new and highly significant.

I do have some comments about the manuscript:

- Existing literature [Beverland et al, JOURNAL OF MATHEMATICAL PHYSICS 57, 022201 (2016)] suggests that the operations that can be carried out transversally, in a 2D topological code, are only those due to braiding abelian anyons, and due to acting with automorphisms of the anyons. The challenge is that the operations needed to move non-abelian anyons (let's call these string operators) are not finite depth unitary circuits. These older works do not take into account the possibility of leveraging measurements, however. The authors should clarify in the main text what their claim about the gate set, carried out using the specific operations that they use, is.

- Experimentally, it would be interesting to hear how much the fact that this particular experiment had to use two qubits, rather than a native qutrit, to realize this model affects fidelity. Would someone who had an architecture with native qutrits have a significant advantage in creating and manipulating this kind of system? Or are these errors negligible compared to, say, those of 2 qubit gates?

In terms of readability, the manuscript is at times so concise that many important technical details are difficult or impossible to follow without turning to the supplementary information. This includes:
— The form of the S_A and S_B operators in Eqs 1 and 2. (Here I also note that "charge" in gauge theories often refers to the irreducible representation, whereas here the authors mean a conjugacy class).
— Circa line 158:
"The nontrivial braiding between the two flux types toggles the fusion channel of the C3 pair from the vacuum to C3"
The operation being discussed is between an excitation (flux) and the flux through one of the non-contractible loops of the torus, which labels a ground state sector. This phrasing could be clarified.
— Figures 3 and 4: The measurements indicate that in these states, many of the plaquette terms are far from being in their ground state value, whereas the caption says that the system should be close to its ground state except near quirt endpoints. Potentially the strings connecting the pairs of quirts at each time step could be added to the Figure to clarify.

Referee #2

(Remarks to the Author)
This manuscript demonstrates the preparation and manipulation of non-Abelian anyons associated with the group S_3 on the Quantinuum H2 trapped-ion quantum processor. The authors encode logical qutrits in the fusion space of non-Abelian fluxes. This encoding is somewhat non-trivial, as part of the information is stored in the internal space of the two-dimensional charge sector. This differs from the more common fusion-channel encoding, in which logical states are associated directly with anyon types. The authors show that combining braiding operations with anyon fusion measurements enables a universal set of logical gates. They further demonstrate the preparation of a magic state using these primitives.

I find the results to be both technically impressive and conceptually important. Topological quantum computation traditionally relies on the braiding statistics of non-Abelian anyons to implement logical gates. The authors show that universality can be achieved when anyon fusion measurements are incorporated into the gate set. In particular, the universal gate set consists of the Pull-through gate together with X- and Z-basis measurements. These primitives are experimentally demonstrated through the creation, braiding, and fusion of non-Abelian anyons. Demonstrating this framework on quantum hardware represents a significant advance, both for the experimental study of topological phases and for the development of fault-tolerant quantum computation protocols.

This work should be of broad interest to researchers working on topological phases of matter, quantum computation, and quantum simulation on near-term quantum devices. Overall, the manuscript is well written and clearly structured. The physical picture and computational framework are presented in a manner that should be accessible to the broad readership of Nature, while the technical aspects appear to be carefully validated.
I only have a few minor suggestions that the authors may wish to consider:

1. While this approach to quantum computation using non-Abelian anyons is very interesting, it is not entirely clear how it compares with other approaches in terms of resource cost. It would be helpful if the authors could comment on the potential advantages or trade-offs of performing topological quantum computation with these non-Abelian excitations.

2. In the section “Bureau of standards,” after performing the Z-measurements the authors post-select on the [-] fusion channel to complete the protocol. What is the acceptance rate associated with this post-selection? Does this imply that the relative encoding implemented in this protocol may face scalability limitations?

3. In the Supplementary Materials, below Eq. S62, the authors state that the “projectors can be shown to be equivalent to the canonical quantum double projectors.” It would be helpful if the authors could clarify in what sense the two constructions are equivalent. For example, are they equivalent only within the ground-state subspace, or are they related by a finite-depth unitary circuit?

These points are purely clarificatory and do not affect the main conclusions.

In summary, I find the work to be novel, technically sound, and of broad interest. I recommend publication in Nature.

Referee #3

(Remarks to the Author)
This paper studies the use of a non-abelian topological order (TO) build from the S3 group as a method to achieve topologically protected universal quantum computation on gate based quantum computers. It focuses on this particular TO as it occupies a “sweet spot” of being easy to prepare—meaning it can be prepared using a finite depth circuit and mid-circuit measurements—while still capable of universal operations. For the latter they employ both braiding and fusion (charge measurements) to achieve a universal gate set and read-out. The team demonstrates the state preparation and a minimal set of operations on Quantinuum’s trapped ion quantum computer.

This work is not the first to propose using braiding of anyons on gate-based quantum computers as a method of achieving universal fault-tolerant computation, nor is it the first to demonstrate the necessary building blocks in experiment. As far as I can tell, the main novelties are in the “efficient” state preparation and the use of non-standard methods for achieving universality that are necessitated by the chosen topological order. The latter notably include the use of charge measurements to supplement braiding, as well as the use of the two-dimensional [2] charge requiring a bureau of standards. The experimental results are certainly impressive, and of the highest scale and quality of any results currently available. However, there is little to no discussion of how the operations that are being performed are actually topologically protected. Before I can recommend this paper for publication in Nature, I would need the authors to address my following comments, particularly to satisfy the criteria of novelty. I believe there is no doubt in the soundness of the results or that the work is important and relevant to a broad audience.

I believe a crucial aspect of this work is the choice of topological order. My understanding is that this is one of few examples of topological order that are sufficient to achieve universal operations, while the states are possible to prepare with finite depth circuits using mid-circuit measurements. These state preparation techniques are an important contribution from a subset of the authors that has been made in recent years. Without this preparation requirement, there are indeed simpler topological orders (such as the double fibonacci TO) that are universal through braiding alone, which have also been

demonstrated experimental, including a paper cited in this work [Nature Physics 20, 1469–1475 (2024)]. I am singling out that paper in particular as I believe that its aim is very close to this current work. This naturally leads to a few questions:

1) The authors find that the measurement-based preparation leads to lower fidelity than the sequential preparation. Do the authors understand why this is the case? Furthermore, while the measurement-based approach is naively more scalable, can they guarantee that the sources of error observed here won't scale badly?
2) I think the context with existing experiments, such as [Nature Physics 20, 1469–1475 (2024)] are not sufficiently clear. While the approach is different, it is not clear how the two would compare practically.
3) For me the most crucial point that is not addressed here is robustness to noise. Since you are working with states, without a Hamiltonian or active correction, I believe that your operations are not fault-tolerant. Both thermal fluctuations, and coherent errors in the moving of anyons can lead to logical errors that will in principle not be suppressed (may even be amplified) on scaling up. While you cite some work that looks at error correction for non-abelian orders, it is not clear how applicable these results are when your operations rely on excitations and braiding. In fairness, I think this is a point that has not been satisfactorily addressed in the literature, and if I have missed it then it should certainly take more prominence in this text. If you can not verify the robustness, then I find statements like in line 208 "We note that this approach is fully topological…" very misleading!
4) Again comparing to [Nature Physics 20, 1469–1475 (2024)], the initial state preparation is actually a small part of the process. The moving, braiding and measuring of charge will be more costly in practice. Can you comment on how your approach, which uses qutrit-qubit pairs and a more complicated topological order with 8 anyon types, compares in practice to the qubit-based double fibonnaci model with 4 anyon types?

Coming back to the novelty of this work. The final sentence of your abstract says "This work establishes S3 TO as simple enough to be prepared efficiently, yet rich enough to enable universal topological quantum computation.", however, I would argue that was done in a previous work by some of the current authors [arXiv:2502.14974], as well as another work [npj Quantum Inf 11, 112 (2025)] that was posted on the arXiv in 2024. Your work certainly includes experimental demonstration of the key components, which are very impressive in themselves. However, this demonstration appears to be the sole source of novelty, and crucial discussion of fault-tolerance of this approach is missing.

I believe that the contents of the submitted work are sound, and when combined with the supplemental material are sufficiently detailed. Supporting data is also freely hosted on Zendo, which is great. I do feel that the details of your implementation are a bit light in the main text, particularly regarding the coherent moving and braiding of anyons. While a lot of detail is contained in the supplemental material, I believe that this should be more transparent. This aspect is crucial for verifying the fault-tolerance of your operations. The experiments performed are very impressive, and data is presented very clearly.

Overall, I think this is a very well-written and important piece of work. The experiments performed are really state of the art. I also think that, despite my comments on novelty, that it is important to consider non-standard approaches to fault tolerant quantum computing. Combined with previous works, this paper is an important contribution in this area. My only real hesitation in recommending this work for publication in Nature is its novelty, particularly in the context of previous work by some of the submitting authors, and in the absence of a discussion of topological protection from thermal and coherent noise. Because of the impressive combination of techniques and ideas that are being demonstrated together, and the string of genuinely novel and powerful results from the current authors that have led to this work, if you are able to address my above comments and reframe your contribution I would likely be able to recommend this work for publication in Nature.

Response to Referees

Response to Referee 1

> *This paper describes experimental realization of the non-abelian topological order associated with the group S3. This is significant because, as a topological phase of matter, S3 is universal for quantum computation — meaning that operations that move and fuse its quasiparticles can be used to carry out a universal gate set. The experimental data presented indicates impressive fidelity with the actual topological state, both after initial state preparation and after elementary computationally relevant operations needed to comprise a universal gate set.*
>
> *Overall I think this work represents a significant advance, worthy of publication in Nature. In particular, while it is not the first experimental demonstration of preparing non-abelian topological order, the demonstration of a universal set of operations is both new and highly significant.*
>
> *I do have some comments about the manuscript:*

We thank the referee for their positive assessment of the work. We are glad the referee finds our results significant and worthy of publication in *Nature*. We address each comment below.

> *- Existing literature [Beverland et al, JOURNAL OF MATHEMATICAL PHYSICS 57, 022201 (2016)] suggests that the operations that can be carried out transversally, in a 2D topological code, are only those due to braiding abelian anyons, and due to acting with automorphisms of the anyons. The challenge is that the operations needed to move non-abelian anyons (let's call these string operators) are not finite depth unitary circuits. These older works do not take into account the possibility of leveraging measurements, however. The authors should clarify in the main text what their claim about the gate set, carried out using the specific operations that they use, is.*

Response. We thank the referee for raising this point. The Beverland et al. result applies specifically to the setting of transversal gates on logical states encoded in the ground state subspace, in the absence of measurements, which is indeed a distinct setting from the one considered in our work.

Our gate set is not claimed to be transversal in the usual sense of unitary operations. Rather, it consists of three explicitly non-transversal primitives which are, however, implementable using constant-depth adaptive circuits (using measurement and feedforward). Recent theory work has argued that this feature can be used to make such an S_3 code fault-tolerant [arXiv:2602.11258].

The three non-transversal (but still can be made fault-tolerant) primitives are: (i) coherent anyon braiding, (ii) X-basis measurement via anyon fusion, and (iii) Z-basis measurement via anyon fusion. The universality claim — that any quantum operation on qutrits (and qubits) can be synthesized from these three primitives, which exploits the cyclic fusion rules of the S_3 quantum double together with mid-circuit anyon fusion measurements. This kind of operation is not considered in Beverland et al.

We have added a clarifying remark in the main text (see line 84-88):

"While these primitives cannot be realized in a strictly transversal way [50, 51], and we use linear-depth circuits with measurements for our demonstrations, they can be scalably implemented with constant-depth adaptive circuit."

distinguishing our measurement-assisted approach from the transversal gate setting of Beverland et al., and making precise which operations constitute our claimed universal gate set. **We have**

also added a reference to the theory work [arXiv:2602.11258] which explains how the constant-depth adaptive circuits enable fault-tolerance.

> *- Experimentally, it would be interesting to hear how much the fact that this particular experiment had to use two qubits, rather than a native qutrit, to realize this model affects fidelity. Would someone who had an architecture with native qutrits have a significant advantage in creating and manipulating this kind of system? Or are these errors negligible compared to, say, those of 2 qubit gates?*

Response. We acknowledge the referee for the insightful question. As the referee correctly pointed out, each physical qutrits encoded by two physical qubits. This encoding introduces a leakage error channel: errors that take the qubit pair outside the computational qutrit subspace $\{|00\rangle, |01\rangle, |10\rangle\}$ (encoding qutrit states $|0\rangle$, $|1\rangle$, $|2\rangle$) are detectable, such that we can use post-selection to discard states outside the computational qutrit subspace. We exploit this as an independent heralding condition (described in Section S4 of the Supplementary Information), and approximately 24% of shots are discarded under all heralding criteria combined. **We have added a brief discussion of the heralding in Section II of the main text (see line 100-103):**

"Errors caused by leakage into the extra qubit state can be mitigated through heralding and post-selection (Supplementary Information Section S4)."

referring readers to Section S4 in Supplementary Information for the detailed error analysis.

An architecture with native qutrits would reduce circuit depth and likely improve fidelity, particularly for the deeper circuits used in the pull-through gate and magic state preparation. However, the magnitude of this advantage is hardware-dependent and would require a direct comparison beyond the scope of this work. For instance, we note that in 2508.16294, the authors numerically compare a native qutrit vs. two-qubit scheme for preparing the $\mathbb{Z}_3$ toric code on a Rydberg atom based platform, and find similar estimated fidelities.

> *In terms of readability, the manuscript is at times so concise that many important technical details are difficult or impossible to follow without turning to the Supplementary Information. This includes: — The form of the S_A and S_B operators in Eqs 1 and 2. (Here I also note that "charge" in gauge theories often refers to the irreducible representation, whereas here the authors mean a conjugacy class).*

Response. We agree that the main text is concise on this point. Figure 1(e–h) provides diagrammatic definitions of all four stabilizers $S_A^{Z_2}$, $S_A^{Z_3}$, $S_B^{Z_2}$, and $S_B^{Z_3}$, and Section S1B gives their explicit Pauli operator forms. **We have elaborated on the inline description of the operator structure to the main text to improve readability (see line 108-114):**

"$S_A^{\mathbb{Z}_2}$ ($S_A^{\mathbb{Z}_3}$) represent the $\mathbb{Z}_2$ ($\mathbb{Z}_3$) X-type stabilizers imposing Gauss law constraints at a vertex, such that the corresponding projector $A_v^{\mathbb{Z}_2}$ ($A_v^{\mathbb{Z}_3}$) projects into the zero-charge sector at vertex v. The plaquette stabilizer $S_B^{\mathbb{Z}_2}$ ($S_B^{\mathbb{Z}_3}$) are Z-type stabilizers detecting $\mathbb{Z}_2$ ($\mathbb{Z}_3$) flux, such that the projector $B_p^{\mathbb{Z}_2} B_p^{\mathbb{Z}_3}$ projects into the trivial-flux sector at plaquette p."

We agree with the referee that the standard usage in gauge theory context for "charge" is to refer to irreducible representation of the gauge group, whereas pure flux refers to conjugacy class. This is consistent with our usage of the term in the main text. However, there may be some confusion with the term "topological charge", which can mean the general anyon type. **For clarity, we have replaced references to "topological charge" with "anyon type".**

— Circa line 158: "The nontrivial braiding between the two flux types toggles the fusion channel of the C3 pair from the vacuum to C3" The operation being discussed is between an excitation (flux) and the flux through one of the non-contractible loops of the torus, which labels a ground state sector. This phrasing could be clarified.

Response. We agree that our phrasing can be made more precise. We start from a ground state in the sector with a C_2 flux threading the vertical non-contractible cycle of the torus (i.e., the ground state is labeled by the topological sector corresponding to a C_2 flux through the vertical loop — this is a ground state, not an excitation). We then create a pair of C_3 flux anyons and braid one of them around the horizontal non-contractible cycle. The non-trivial mutual statistics between a C_3 anyon and the background C_2 topological sector (the flux through the vertical cycle) causes the C_3 pair's total fusion channel to change from vacuum (the initial, pair-created state) to C_3, leaving a net single C_3 anyon in the system. **We have re-written this sentence (see line 174-181) to clarify the distinction between the background topological sector (a property of the ground state) and a localized excitation:**

"Braiding a C_2 flux pair around the vertical noncontractible cycle of the torus generates a nontrivial ground state with C_2 background flux. Starting from this nontrivial ground state, we can braid a C_3 flux pair around the horizontal noncontractible cycle (Fig. 2b). Due to the nontrivial braiding between the background C_2 flux and the C_3 anyons, the C_3 pair fuses to a single C_3 anyon after braiding around the horizontal cycle (Fig. 2c)."

— Figures 3 and 4: The measurements indicate that in these states, many of the plaquette terms are far from being in their ground state value, whereas the caption says that the system should be close to its ground state except near quirt endpoints. Potentially the strings connecting the pairs of quirts at each time step could be added to the Figure to clarify.

Response. We appreciate the referee for the suggestion. In Fig. 3, the $B^{\mathbb{Z}_3}$ plaquette values away from 1 denotes a violation, which is the result of the C_2 flux at the plaquette. The white plaquettes denotes the location of such violations, whereas the green plaquettes are for values of $B^{\mathbb{Z}_3} \approx 1$, i.e. no flux at that plaquette. **To improve clarity for the reader, we have added string connecting the qutrit pair to Fig. 3(c) and the new combined Fig. 4(b).**

Response to Referee 2

> *This manuscript demonstrates the preparation and manipulation of non-Abelian anyons associated with the group S_3 on the Quantinuum H2 trapped-ion quantum processor. The authors encode logical qutrits in the fusion space of non-Abelian fluxes. This encoding is somewhat non-trivial, as part of the information is stored in the internal space of the two-dimensional charge sector. This differs from the more common fusion-channel encoding, in which logical states are associated directly with anyon types. The authors show that combining braiding operations with anyon fusion measurements enables a universal set of logical gates. They further demonstrate the preparation of a magic state using these primitives.*
>
> *I find the results to be both technically impressive and conceptually important. Topological quantum computation traditionally relies on the braiding statistics of non-Abelian anyons to implement logical gates. The authors show that universality can be achieved when anyon fusion measurements are incorporated into the gate set. In particular, the universal gate set consists of the Pull-through gate together with X- and Z-basis measurements. These primitives are experimentally demonstrated through the creation, braiding, and fusion of non-Abelian anyons. Demonstrating this framework on quantum hardware represents a significant advance, both for the experimental study of topological phases and for the development of fault-tolerant quantum computation protocols.*
>
> *This work should be of broad interest to researchers working on topological phases of matter, quantum computation, and quantum simulation on near-term quantum devices. Overall, the manuscript is well written and clearly structured. The physical picture and computational framework are presented in a manner that should be accessible to the broad readership of Nature, while the technical aspects appear to be carefully validated.*
>
> *I only have a few minor suggestions that the authors may wish to consider:*

We thank the referee for their thorough and positive assessment of the clarity and significance of the work. We address each suggestion below.

> *1. While this approach to quantum computation using non-Abelian anyons is very interesting, it is not entirely clear how it compares with other approaches in terms of resource cost. It would be helpful if the authors could comment on the potential advantages or trade-offs of performing topological quantum computation with these non-Abelian excitations.*

Response. This is an important question for situating our work within the broader landscape of topological quantum computation. **We have added discussions of resource costs and trade-offs to various parts of the main text, addressing the following key comparisons:**

Comparison with Fibonacci / braid-only universal approaches: These points of comparison are summarized in the introduction. The Fibonacci topological order achieves universality through braiding alone, which is conceptually elegant. However, this comes at significant practical cost: (a) there is no known constant-depth adaptive preparation protocol for Fibonacci TO, as conjectured no-go's for solvable-group preparation techniques exclude it (Ref. [15] in main text); (b) Fibonacci anyons suffers from complex leakage issues due to the irrational quantum dimension of its anyons, making leakage-free universal gate sets difficult to construct (Refs. [34]–[36] in main text); (c) there are no decoders with a proven threshold theorem for Fibonacci anyons. By contrast, the S_3 quantum double is solvable, admits a constant-depth adaptive preparation (demonstrated in our work), and the integer quantum dimensions of anyons in S3 quantum double encoding enable more

straightforward leakage correction (Ref. [9] in main text). Furthermore, it has recently been proven that the S_3 quantum double has a fault-tolerance threshold (newly included Ref. [37] in main text). **This comparison is adapted in line 39-51 in the Introduction:**

"This stands in contrast to more exotic topological phases such as Fibonacci TO where braid-only universality can be achieved [4], albeit at the cost of significantly greater complexity in their realization [29,31] in part because of conjectured no-go's of scalable preparation protocols [15]. In addition, their irrational quantum dimensions complicate the search of leakage-free gates [34-36], in contrast to the integer quantum dimensions of anyons in quantum doubles. Moreover, while a finite decoding threshold has recently been proven for quantum doubles with solvable groups by leveraging the similarity with the toric code [37], no similar results exist for more exotic TO."

Comparison with surface code / ground-state-based approaches: Conventional topological codes (toric code, surface code) encode information in the ground-state subspace and achieve fault-tolerance via active error correction. Our approach encodes directly in the fusion space of anyonic excitations, which provides an alternative encoding with different trade-offs: the code distance scales with the separation between anyons (tunable in principle), but active error correction for non-Abelian anyons is more complex. The overhead of the bureau of standards and the repeat-until-success requirements (see Comment 2 below) are costs specific to our approach that would need to be weighed against the benefits. While a quantitative comparison would be very interesting, it would naturally depend on various hardware-specific details and would deserve its own body of work. **We have updated the following sentences in the outlook (see line 434-444) to discuss this comparison:**

"Since our gate set is topological, the operations are protected as long as one is in the phase. Stabilizing the topological phase requires active error correction, which is beyond the scope of the present work, although a recent work has proven a finite decoding threshold even with imperfect measurement during braiding [37]. Numerical works indicate thresholds may be comparable to that of the surface code [54,55], and a systematic comparison between ground-state and fusion-space encoding is an exciting open direction."

Resource overhead of our primitives: The pull-through gate requires moving one anyon endpoint around two others; the X- and Z-basis measurements each require ancilla anyon pairs and fusion events. These are all implemented by linear-depth circuits in this work due to the small system size, but they can be implemented via finite-depth adaptive circuits when system size is scaled up. **We added a clarifying remark in the main text (see line 84-88):**

"While these primitives cannot be realized in a strictly transversal way [50, 51], and we use linear-depth circuits with measurements for our demonstrations, they can be scalably implemented with constant-depth adaptive circuit."

> *2. In the section "Bureau of standards," after performing the Z-measurements the authors post-select on the [-] fusion channel to complete the protocol. What is the acceptance rate associated with this post-selection? Does this imply that the relative encoding implemented in this protocol may face scalability limitations?*

Response. We thank the referee for this important clarifying question. The post-selection on the $[-]$ fusion channel in the bureau of standards corresponds to projecting onto the subspace of distinct Z-basis states across the three flux pairs. In the idealized noiseless case, each comparison measurement between uncorrelated flux pair yields the $[-]$ outcome with probability $\frac{1}{2}$. Therefore,

by post-selecting at the end of the three [2] charge comparison measurements, the theoretical value acceptance rate of $(\frac{1}{2})^3 = \frac{1}{8}$. In the experiment discussed in the main text, we post-selected all three comparison measurements on [−] simultaneously at the end of the protocol, and the acceptance rate is ~ 0.115, comparable with the theoretical value.

We want to point out that, as noted in the main text at the end of Section IV.E, post-selection is not necessary: the repeat-until-success strategy — creating fresh C_2 flux pairs in $|\tilde{0}\rangle_L$ and comparing until the [−] outcome is obtained — provides a probabilistic but unbiased preparation of the bureau of standards without any exponential overhead in the number of attempts. Each round of comparison measurement succeeds with a constant probability, so the expected number of rounds is constant and does not scale with system size. Moreover, in practice one would use all measurement outcomes and put the resulting state in the corresponding bin, as a sort of factory.

We have added the explicit post-selection acceptance rate to the Extended Data Figure on the $\mathcal{Z}$ bureau of standards.

> *3. In the Supplementary Information, below Eq. S62, the authors state that the "projectors can be shown to be equivalent to the canonical quantum double projectors." It would be helpful if the authors could clarify in what sense the two constructions are equivalent. For example, are they equivalent only within the ground-state subspace, or are they related by a finite-depth unitary circuit?*

Response. We thank the referee for the clarification question. There are two steps to obtaining the projectors in Eq. S62: 1) the group-valued quantum double projectors are re-expressed in terms of the qutrit-qubit encoding, and 2) the projectors in the qutrit-qubit encoding has a simplification within the ground state subspace.

For the first part, the two constructions are related by a relabeling of local group element degrees of freedom that corresponds to the semidirect product decomposition $S_3 = Z_3 \rtimes Z_2$: the local qutrit-qubit pair faithfully encodes the S_3 group element $g = \mu^a \sigma^b$ (with qutrit in state $|a\rangle$ and qubit in state $|b\rangle$), and the stabilizer operators decompose accordingly into Z_3 and Z_2 components acting on the qutrit and qubit respectively. We provide a derivation of this mapping in Section S2.

For the second part, the $\mathbb{Z}_3$ plaquette projector can be further simplified within the ground state subspace. **We have added a clarification (Supplementary Information line 310-312) to explain how the canonical quantum double projector reduces to the form given in Eq. 62:**

"The resulting projectors can be shown to be equivalent to the canonical quantum double projectors (see Section S2) within the ground state subspace of the $\mathcal{D}(S_3)$ model. In particular, the vertex projectors and the $B^{\mathbb{Z}_2}$ are exactly the same, whereas the $B^{\mathbb{Z}_3}$ given in Eq. S34 in Section S2 can be reduced to the above expression when $B^{\mathbb{Z}_2} = +1$."

> *These points are purely clarificatory and do not affect the main conclusions.*
>
> *In summary, I find the work to be novel, technically sound, and of broad interest. I recommend publication in Nature.*

Response. We thank the referee for the positive assessment and recommendation.

Response to Referee 3

> *This paper studies the use of a non-abelian topological order (TO) build from the S3 group as a method to achieve topologically protected universal quantum computation on gate based quantum computers. It focuses on this particular TO as it occupies a "sweet spot" of being easy to prepare—meaning it can be prepared using a finite depth circuit and mid-circuit measurements—while still capable of universal operations. For the latter they employ both braiding and fusion (charge measurements) to achieve a universal gate set and read-out. The team demonstrates the state preparation and a minimal set of operations on Quantinuum's trapped ion quantum computer.*
>
> *This work is not the first to propose using braiding of anyons on gate-based quantum computers as a method of achieving universal fault-tolerant computation, nor is it the first to demonstrate the necessary building blocks in experiment. As far as I can tell, the main novelties are in the "efficient" state preparation and the use of non-standard methods for achieving universality that are necessitated by the chosen topological order. The latter notably include the use of charge measurements to supplement braiding, as well as the use of the two-dimensional [2] charge requiring a bureau of standards. The experimental results are certainly impressive, and of the highest scale and quality of any results currently available. However, there is little to no discussion of how the operations that are being performed are actually topologically protected. Before I can recommend this paper for publication in Nature, I would need the authors to address my following comments, particularly to satisfy the criteria of novelty. I believe there is no doubt in the soundness of the results or that the work is important and relevant to a broad audience.*

Response. We thank the referee for their detailed and thoughtful critique, and for their generous assessment of the experimental quality and the importance of the work. The referee raises interesting points that inspired revisions which we believe has strengthened our manuscript.

We will elaborate on the novelty of our work in our following detailed response.

> *I believe a crucial aspect of this work is the choice of topological order. My understanding is that this is one of few examples of topological order that are sufficient to achieve universal operations, while the states are possible to prepare with finite depth circuits using mid-circuit measurements. These state preparation techniques are an important contribution from a subset of the authors that has been made in recent years. Without this preparation requirement, there are indeed simpler topological orders (such as the double fibonacci TO) that are universal through braiding alone, which have also been demonstrated experimental, including a paper cited in this work [Nature Physics 20, 1469–1475 (2024)]. I am singling out that paper in particular as I believe that its aim is very close to this current work. This naturally leads to a few questions:*
>
> *1) The authors find that the measurement-based preparation leads to lower fidelity than the sequential preparation. Do the authors understand why this is the case? Furthermore, while the measurement-based approach is naively more scalable, can they guarantee that the sources of error observed here won't scale badly?*

Response. We thank the referee for the incisive question. The overall fidelity in the Quantinuum's H2 hardware is upper-bounded by the fidelity of the native two-qubit gate (around 0.998), which is the dominant source of error. Therefore, a quick back-of-the-envelope estimation of the ground state preparation fidelity is to use the following formula

$$\text{(Fidelity per edge)} = \text{(Fidelity of two-qubit gate)}^{\frac{d}{N}} \tag{1}$$

where d is the total number of native two-qubit gates used in the preparation circuit, and $N = 3 \times 3 \times 2 = 18$ is the number of edges of the square lattice (for our 3×3 geometry).

For the unitary preparation protocol, the circuit depth is reported in the Supplementary Table S3, for the ground state unitary preparation, the native two-qubit gate depth $d = 238$. Therefore, the fidelity per edge is $(0.998)^{238/18} \approx 0.974$, which is within the fidelity bound of 0.970(5)–0.988(3) given in Eq. 3 in the main text.

For the measurement-based adaptive preparation protocol, the native two-qubit gate depth $d = 274$. Therefore, the fidelity per edge is $(0.998)^{274/18} \approx 0.970$, which is within the fidelity bound of 0.930(8)–0.978(2) given right below Eq. 3 in the main text. The fidelity is slightly worse than unitary preparation protocol because of the constant overhead involved for measurement-based protocol. Indeed, the measurement-based protocol requires slightly more gates since we need to entangle to ancillas (which we subsequently measure).

For the current system sizes, error contributions are dominated by gate count rather than runtime; therefore, the unitary preparation has (slightly) better fidelity. Naturally, as system size is scaled up, the runtime of the unitary preparation scales linearly with system size and will eventually form a bottleneck, such that the adaptive protocol then wins out. More precisely, for the adaptive protocol, the number of gates grows proportional to system size ($\sim L^2$), while the preparation time remains constant, such that the fidelity-per-qubit will remain constant (which can then be combined with quantum error correction if this local fidelity is beyond some stability threshold). In contrast, for the unitary preparation, while the number of gates also grows as $\sim L^2$, the preparation time scales as $\sim L$, leading to a fidelity $\sim e^{-aL^3}$, which means $e^{-aL} \sim e^{-a\sqrt{N}}$ per qubit.

We thank the referee for raising this comparison; we have augmented Extended Data Table to include the gate counts for the measurement-based protocol to aid comparison. We have also added the above back-of-the-envelope estimation and scaling asymptotics to the Supplementary Information Section 6C on the ground state preparation.

> *2) I think the context with existing experiments, such as [Nature Physics 20, 1469–1475 (2024)] are not sufficiently clear. While the approach is different, it is not clear how the two would compare practically.*

Response. We agree that a more direct and explicit comparison is warranted. The Fibonacci anyon experiment of Xu et al. [*Nature* Physics 20, 1469 (2024)] realized braid-only universality in a superconducting processor using non-Abelian Fibonacci anyons. The two approaches differ practically in several crucial respects:

Scalable state preparation: Fibonacci TO has no known constant-depth adaptive preparation protocol, and it is conjectured (Ref. [15] in main text) that no such protocol exists for Fibonacci TO. S_3 is solvable, and the constant-depth adaptive preparation demonstrated in our work is a key practical advantage for scaling.

Gate leakage: Fibonacci anyons require careful management of leakage out of the computational subspace because of the irrational quantum dimension of Fibonacci anyons; constructing leakage-free universal gate sets is a known challenge (Refs. [34]–[36] in main text). Existing experiments, including Nature Physics 20, 1469–1475 (2024)], do not address the gate leakage issue. In

contrast, the integer quantum dimensions of anyons in S_3 quantum double encoding enable more straightforward leakage correction (Ref. [9] in main text).

Error correction: To our knowledge, there is no proof of a finite code threshold for Fibonacci anyons; there are only numerical results that suggest a finite threshold [Phys. Rev. A 95, 022309 (2017), Phys. Rev. X 12, 021012 (2022)]. On the other hand, there is a recent theory proof for a finite code threshold of solvable anyon theory (of which S_3 quantum double is an example) using a just-in-time decoder (newly included Ref. [37] in main text). We also note that a recent numerical work that suggests that such non-Pauli stabilizer code using just-in-time decoder has a threshold comparable to standard global decoder (newly included Ref. [55] in main text). Therefore, with the established code threshold, computing with S_3 quantum double can be made fault tolerant by scaling up system size.

We revised the Introduction (see line 39-51) to include a clearer, more substantive comparison:

"This stands in contrast to more exotic topological phases such as Fibonacci TO where braid-only universality can be achieved [4], albeit at the cost of significantly greater complexity in their realization [29,31] in part because of conjectured no-go's of scalable preparation protocols [15]. In addition, their irrational quantum dimensions complicate the search of leakage-free gates [34-36], in contrast to the integer quantum dimensions of anyons in quantum doubles. Moreover, while a finite decoding threshold has recently been proven for quantum doubles with solvable groups by leveraging the similarity with the toric code [37], no similar results exist for more exotic TO."

> *3) For me the most crucial point that is not addressed here is robustness to noise. Since you are working with states, without a Hamiltonian or active correction, I believe that your operations are not fault-tolerant. Both thermal fluctuations, and coherent errors in the moving of anyons can lead to logical errors that will in principle not be suppressed (may even be amplified) on scaling up. While you cite some work that looks at error correction for non-abelian orders, it is not clear how applicable these results are when your operations rely on excitations and braiding. In fairness, I think this is a point that has not been satisfactorily addressed in the literature, and if I have missed it then it should certainly take more prominence in this text. If you can not verify the robustness, then I find statements like in line 208 "We note that this approach is fully topological..." very misleading!*

Response.

The referee raises a critical and valid point that we should address head-on. The statement that "this approach is fully topological" refers specifically to the logical encoding and gate operations: the encoded information is stored non-locally in the fusion space of separated anyons, and the gates are implemented through braiding and fusion — operations that are topological in the sense of depending only on the topology of anyon worldlines, not on their precise spatial paths. **We edited the sentence (line 230-234) to "We note that this approach is fully topological, in the sense that no individual local measurement can corrupt the logical information; and if so desired, it can be expressed in terms of fusion trees (Supplementary Information Section S9)" for clarity.**

This is separate from the question of how to stabilize the topological phase of matter in and of itself. We acknowledge we do not address this, and in practice it can be achieved either in solid state state systems via energy gaps, or in quantum platforms via active error correction. **We added a discussion in the outlook (see line 434-444) to reflect this:**

"Since our gate set is topological, the operations are protected as long as one is in the phase. Stabilizing the topological phase requires active error correction, which is beyond the scope of the present work, although a recent work has proven a finite decoding threshold even with imperfect measurement during braiding [37]. Numerical works indicate thresholds may be comparable to that of the surface code [54,55], and a systematic comparison between ground-state and fusion-space encoding is an exciting open direction."

> *4) Again comparing to [Nature Physics 20, 1469–1475 (2024)], the initial state preparation is actually a small part of the process. The moving, braiding and measuring of charge will be more costly in practice. Can you comment on how your approach, which uses qutrit-qubit pairs and a more complicated topological order with 8 anyon types, compares in practice to the qubit-based double fibonnaci model with 4 anyon types?*

Response.

We appreciate the referee for the question. We believe that the number of anyons is not a direct measure of complexity for quantum computation: while S_3 topological order has more anyon types than Fibonacci anyons, we are only using one anyon type, C_2 anyon pair, to encode a logical qutrit. S_3 topological order is considerably simpler than Fibonacci anyons in various ways, for instance leakage-free gate (since S_3 anyons have integer quantum dimensions), 2-anyon encoding, and simple compilations of gates (whereas Fibonacci suffers from complicated compilation for even some Clifford gates). Moreover, the fact a fault-tolerance threshold has been proven for S_3 topological order (newly included Ref. [37] in main text), but not for Fibonacci anyons, is another indication that S3 topological order is simpler. This is related to the fact that although S3 has more anyons which all enter the error correction protocol, their fusion and braiding properties are more manageable than those of Fibonacci (which is indeed responsible for S3 not being braiding-universal) which in turn makes the error correction less challenging.

Let us also comment on the initial state preparation: while this is indeed a small part, it is very advantageous to immediately start very close to the idealized topological order, since quantum error correction is only guaranteed to work when one is sufficiently close to the desired state. Intuitively, the existence of the measurement-based state preparation enables one to immediately 'turn on' the error correction protocol after the state preparation. In contrast, it is currently not clear to us how to prepare the Fibonacci topological order in a way that one can be guaranteed to be below the error correction threshold (since a sufficiently large system requires a sufficiently deep circuit, which means the local fidelity may quickly deteriorate as discussed in response to question #1 above).

> *Coming back to the novelty of this work. The final sentence of your abstract says "This work establishes S3 TO as simple enough to be prepared efficiently, yet rich enough to enable universal topological quantum computation.", however, I would argue that was done in a previous work by some of the current authors [arXiv:2502.14974], as well as another work [npj Quantum Inf 11, 112 (2025)] that was posted on the arXiv in 2024. Your work certainly includes experimental demonstration of the key components, which are very impressive in themselves. However, this demonstration appears to be the sole source of novelty, and crucial discussion of fault-tolerance of this approach is missing.*

Response.

We thank the referee for the inquiry. We agree that our initial manuscript undersold some of our theoretical innovation, and we thank the referee for motivating us to make this more explicit. In addition to the experimental demonstrations of the universal gate set and the experimental realization of a state with single *non-Abelian* anyon on a torus, the primary novel theoretical contributions of the present work are:

1. First two-anyon encoding: Our work presents an encoding based on only a pair of non-Abelian flux, since we make use of the bureau of standards. This is more efficient than our previous 4-anyon encoding we presented in our prior work [arXiv:2502.14974] and in [npj Quantum Information 11, 112 (2025)], which do not make use of the bureau of standards.

2. Coherent Moving: We developed unitary circuits for moving non-Abelian fluxes in quantum double using conditioned ribbon operators without altering the logical state. While there are a variety of works writing down explicit ribbon operators for anyons, we are not aware of these protocols for coherently moving anyons leveraging the ribbon operators in quantum double models.

3. Fusion tree equivalence: We showed that our logical encoding based on bureau of standards can be made fully consistent with the fusion-tree understanding of topological quantum computation. We explicitly give an example of our logical state in the fusion tree basis in Supplementary Information S9. We note that in previous works, these two approaches have often been treated as conceptually different [Quantum Information Processing 14: 2687-2727 (2015)].

   We added the following discussion in the outlook of the main text (see line 408-415) to highlight our aforementioned theoretical contributions:

   "These experimental demonstrations are also enabled by theoretical advances: (i) we showed how a two-anyon encoding is sufficient in contrast to proposed four-anyon encodings [46,47], (ii) we developed a coherent moving protocol to move non-Abelian flux deterministically, and (iii) we showed that the logical encoding based on bureau of standards can be described in the fusion-tree perspective."

   We acknowledge the referee's comment that a few aspects of the theoretical case for S_3 TO as a platform for universal topological quantum computation has been made in prior works, including [arXiv:2502.14974] and [npj Quantum Information 11, 112 (2025)]. **We have added a reference to [npj Quantum Information 11, 112 (2025)] and revised the abstract sentence to "This demonstrates S_3 TO to be scalably preparable yet rich enough to support a universal gate set" to capture both the theoretical and experimental advances of our work.**

> *I believe that the contents of the submitted work are sound, and when combined with the Supplementary Information are sufficiently detailed. Supporting data is also freely hosted on Zendo, which is great. I do feel that the details of your implementation are a bit light in the main text, particularly regarding the coherent moving and braiding of anyons. While a lot of detail is contained in the Supplementary Information, I believe that this should be more transparent. This aspect is crucial for verifying the fault-tolerance of your operations. The experiments performed are very impressive, and data is presented very clearly.*

Response. We agree with the referee that the main text is currently sparse on the implementation details of coherent anyon transport. **We expanded the main text description of the coherent anyon moving protocols (see line 239-242):**

"The braiding of non-Abelian fluxes is enabled by technical advancements in protocols to coherently move non-Abelian anyons using local flux label while preserving the global logical information [53] (see Supplementary Information Section S3 for theory overview)."

We also added a statement in Supplementary Information Section S3 (see line 177-179) clarifying the key requirement for topological protection:

"We also require that, at no point in the protocol, is the global flux label (i.e. what defines the "absolute" logical encoding) stored locally; the coherent moving protocol only involves the local flux label instead."

> *Overall, I think this is a very well-written and important piece of work. The experiments performed are really state of the art. I also think that, despite my comments on novelty, that it is important to consider non-standard approaches to fault tolerant quantum computing. Combined with previous works, this paper is an important contribution in this area. My only real hesitation in recommending this work for publication in Nature is its novelty, particularly in the context of previous work by some of the submitting authors, and in the absence of a discussion of topological protection from thermal and coherent noise. Because of the impressive combination of techniques and ideas that are being demonstrated together, and the string of genuinely novel and powerful results from the current authors that have led to this work, if you are able to address my above comments and reframe your contribution I would likely be able to recommend this work for publication in Nature.*

Response. We thank the referee for the careful and overall positive assessment of our manuscript. We hope that in our responses above, we have sufficiently addressed the referee's concern on the novelty and framing of our work.